

A Bayesian Framework for Landing Site Selection during Autonomous Spacecraft Descent^{*}

Navid Serrano

Jet Propulsion Laboratory

California Institute of Technology

4800 Oak Grove Dr, Pasadena, CA 91109, USA

navid.serrano@jpl.nasa.gov

Abstract – The success of a landed space exploration mission depends largely on the final landing site. Factors influencing site selection include safety, fuel-consumption, and scientific return. This paper addresses the problem of selecting the best available landing site based on these factors in real-time during autonomous spacecraft descent onto a planetary surface. The problem is modeled probabilistically using Bayesian Networks (BNs). BNs provide a means of representing the causal relationships between variables that impact the quality of a landing site. The final landing site is determined via probabilistic reasoning based on terrain safety derived from on-board sensors, available fuel based on spacecraft descent dynamics, and regions of interest defined by mission scientists.

Index Terms – *Autonomous spacecraft, safe landing, terrain characterization, Bayesian Networks.*

I. INTRODUCTION

Landing sites for space exploration missions have historically been determined off-line by scientists and engineers based on aerial imagery obtained from orbiters. The selection process involves a variety of criteria including safety, engineering, and science factors [1]. For instance, in the case of the *Mars Exploration Rovers* (MER) mission, as many as 185 possible landing sites were first identified before being narrowed down to 6 high priority science sites. Among the concerns that factored into the selection process were horizontal winds and wind shear, slopes, and rocks on the surface [2]. In the end, two landing sites were selected—one for each rover.

Although, scientific return is a fundamental part of any space exploration mission, ultimately, the dominating concern is safety. Consequently, landing sites with high scientific potential are often eliminated from consideration if safety cannot be guaranteed. For this reason, a major effort is underway to equip the next generation of unmanned spacecraft with onboard hazard detection and avoidance capabilities in order to reach locations of higher scientific interest while meeting the necessary safety criteria.

Considerable work has been done in the area of sensor-based autonomous landing. Landmark detection is of particular importance for spacecraft navigation [3] and landing [4]. In

some cases, known landing hazards, such as craters [4,5], are used as landmarks for both position estimation and safe landing. In other cases, terrain features, such as slope and roughness, are used to determine landing safety [6]. Prior research has generally focused on the use of a single sensor, such as a camera [7] or LIDAR [8]. The work presented here focuses first, on the use of multiple heterogeneous on-board sensors and second, on the use of reasoning techniques to infer safety and incorporate engineering and science factors in the selection of a final landing site.

The use of a reasoning engine for terrain characterization using multiple sensors was first proposed in [6]. Terrain features extracted from a RADAR, LIDAR, and camera were mapped to a multi-level safety scale using fuzzy logic. In this paper, the problem is modelled probabilistically using Bayesian Networks (BNs). BNs not only provide a framework for terrain safety assessment, but also for the selection of an optimal landing site based on additional critical factors, which has not been formally addressed before.

BNs have many attractive characteristics that make them particularly well suited to this problem. Because BNs model the relationships between causes and effects, they can be used for both inference and causal reasoning. In a BN, the safety of the terrain can be inferred from features extracted from the on-board sensors. In addition, by incorporating other factors, such as fuel consumption and scientific return, a BN can also be used to determine the best landing site using causal reasoning. In this paper, a Bayesian framework for landing site selection is presented and discussed. In addition, experiments are conducted using synthetic planetary terrains in order to simulate the landing site selection during spacecraft descent.

II. TERRAIN SAFETY

A. Onboard Sensors

A combination of active and passive sensors is used for terrain characterization during descent: RADAR, LIDAR, and camera, as in [6,9]. The motivation for using multiple sensors is twofold: 1) to increase robustness and 2) to enrich the feature set.

Each sensor has different physical characteristics, such as field of view, resolution, and operating range. The limitations

^{*} The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology under contract from the National Aeronautics and Space Administration.

of a particular sensor may be compensated by the strengths of another. In addition, since the sensors have different operating ranges, the fusion of sensory information is grouped into three tiers, as shown in Table I.

TABLE I
TIERED SENSOR OPERATION

Tier	Range	Operational Sensor(s)
1	10km – 7km	Radar
2	7km – 1km	Radar + Camera
3	1km - Touchdown	Radar + Camera + Lidar

B. Terrain Features

A combination of large scale and small scale topographic features are used to predict terrain safety. The small scale features are obtained from RADAR and LIDAR range data. Slope and roughness features can be extracted from the range data using plane fitting.

Let $z = ax + by + c$ represent a plane in \mathfrak{R}^3 . The plane parameters a , b , and c are estimated from the range data using the Least Median of Squares (LMedSq) regression technique [8]. The slope is obtained by calculating the angle formed by the estimated plane normal and the z -axis:

$$f_{\theta}(x, y) = \cos^{-1} \left(\frac{1}{\sqrt{a^2 + b^2 + 1}} \right) \quad (1)$$

where a and b are the parameters of the best-fitting plane at location (x, y) . The fitting error between the plane and the sensor range provides a measure of the local terrain roughness:

$$f_e(x, y) = |d(x, y) - (ax + by + c)| \quad (2)$$

A novel aspect of this work is that in addition to local terrain features, known landing hazards such as craters and rocks are explicitly detected and used in the reasoning process. These are obtained from camera imagery using detection algorithms. The crater detection algorithm uses edges and shadow patterns to identify candidate craters [4]. The candidate craters are parameterized by fitting an ellipse to the crater boundary. Let $x_{0,i}$, $y_{0,i}$, a_i , b_i , and ϕ_i be the ellipse center x -coordinate, center y -coordinate, semi-major axis length, semi-minor axis length, and rotation angle, respectively, for the i th detected crater. The presence of craters is defined as:

$$f_c(x, y) = \begin{cases} 1, & \text{for } \left(\frac{x^2}{a_i^2} + \frac{y^2}{b_i^2} \right) \leq 1 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where x and y are points in a coordinate system rotated by ϕ_i and translated by $x_{0,i}$ and $y_{0,i}$.

At lower altitudes, rocks and boulders are visible and are detected using the algorithm described in [10]. Rocks are detected in descent imagery by locating and characterizing shadows. The shape of the rock is determined using a hemispherical model and the projection of shadows based on the known sun angle. The presence of rocks is defined as:

$$f_r(x, y) = \begin{cases} 1, & \text{for } (x, y) \in R \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where R is the set of pixel locations in the image identified as rocks. Crater and rock detection results are shown in Figure 1.

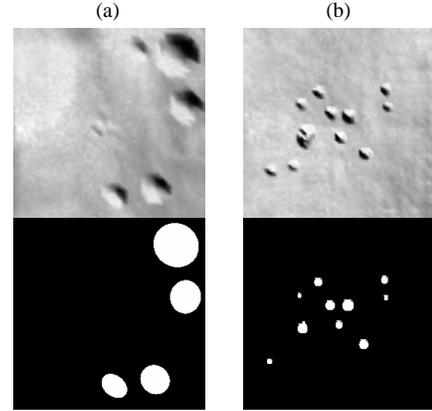


Fig. 1 Craters (a) and rocks (b) detected from camera imagery.

C. Bayesian Safety Assessment

Predicting terrain safety from a set of noisy sensor measurements is not a deterministic problem. For this reason, a probabilistic framework is proposed here. Specifically, Bayesian Networks (BNs) [11] are used to model statistical dependencies and infer terrain safety from features.

BNs are directed acyclic graphs (DAG) where the nodes represent variables and the links between nodes represent causal dependence. The direction of a link indicates causality, and thus a dependence relationship. Nodes that exist at the same level are considered conditionally independent.

The state of the terrain T can be inferred from the terrain features \mathbf{f} using Bayes' rule:

$$P(T | \mathbf{f}) = \frac{P(T, \mathbf{f})}{P(\mathbf{f})} = \frac{P(\mathbf{f} | T)P(T)}{P(\mathbf{f})} \quad (5)$$

If the N features are fully dependent, then an N -dimensional distribution is needed to evaluate Eq. (5). If the features are assumed to be conditionally independent, the N -dimensional distribution is reduced to N 1-dimensional distributions (naïve Bayes). Although independence is often difficult to assess, it has been shown that for most classification problems, the assumption is adequate and does not lead to increased error [12]. Applying conditional independence reduces Eq. (5) to:

$$P(T_k | \mathbf{f}) = \frac{\prod_{i=1}^N P(f_i | T_k) P(T_k)}{\sum_{j=1}^{|T|} \prod_{i=1}^N P(f_i | T_j)} \quad (6)$$

The posterior probability $P(T_k | \mathbf{f}) \in [0, 1]$ provides a continuous-valued measure of certainty that the terrain is safe. Eq. (6) can be represented graphically, as shown in Figure 2.

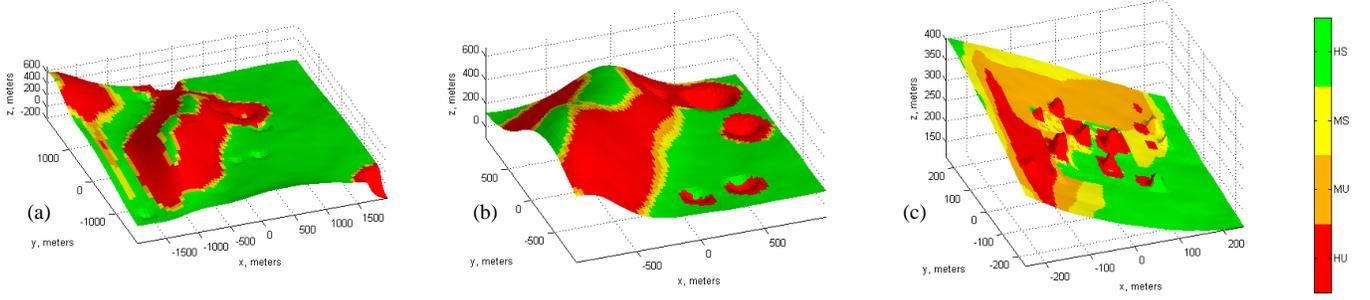


Fig. 3 Terrain safety assessment during descent onto a planetary surface at 8km (a), 4km (b), and 1km (c).

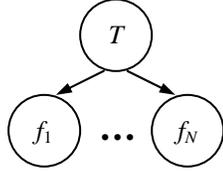


Fig. 2 Naïve Bayes graphical structure for terrain safety assessment.

Example results using Bayesian terrain safety assessment are shown in Figure 3. The terrain safety is shown at three altitudes during descent (onto the same terrain). For display purposes, the continuous probability $P(T|\mathbf{f})$ is mapped to four levels of safety: *highly-unsafe* (HU), *moderately-unsafe* (MU), *moderately-safe* (MS), and *highly-safe* (HS), which are shown as green, yellow, orange, and red, respectively. (The same color scheme is used throughout the paper.)

IV. ENGINEERING FACTORS

Although terrain safety is of paramount importance in landing site selection, other factors must also be considered. During autonomous spacecraft descent, retargeting operations are performed in order to avoid landing hazards and reach the nominal landing site. However, the reachable terrain is constrained by the spacecraft's descent trajectory, velocity and available fuel. Using ballistic analysis, it was shown in [13] that the reachable terrain (landing footprint) is bounded by an ellipse

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1 \quad (7)$$

whose semi-major axis a and semi-minor axis b are defined by

$$a = \frac{(\Delta V^2 - 2E/m)\Delta t}{2\Delta V(1 - v_H/\Delta V)} \quad (8)$$

$$b = a\sqrt{1 - v_H/\Delta V} \quad (9)$$

where ΔV is the allowable change in velocity based on fuel allocation, Δt is the time to impact, v_H is the horizontal velocity, m is spacecraft mass, and E is energy, [13]. It goes

without saying that the landing site must be within the reachable boundary.

Figure 4 shows estimated landing footprints at various points during descent. As shown, the ballistic trajectory begins with an initial horizontal velocity. The landing ellipse then changes dramatically after a re-targeting maneuver is applied. The position of the spacecraft during descent is shown as a red circle and the corresponding landing ellipse is shown in blue.

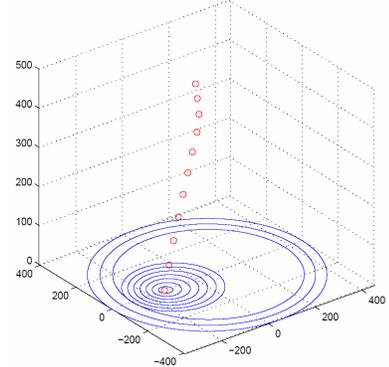


Fig. 4 Estimated landing ellipse using ballistic analysis.

In the Bayesian framework, the landing ellipse is used to exclude regions of the terrain that are unreachable. Let F be a random variable indicating fuel sufficiency. In the binary case, the fuel is either sufficient or insufficient to reach a particular point on the terrain. However, such a representation may not be adequate since the estimate of the landing footprint (and thus fuel sufficiency) is subject to error. One alternative is to shrink the landing ellipse based on the margin of error. Another alternative is to use a multi-level representation. For example, F can be modeled as a ternary discrete random variable where in addition to *reachable* and *unreachable*, a third case is introduced: *marginally reachable*. These are points that lie inside the landing ellipse but are close to the boundary and thus, within the margin of error. Such an approach is adopted here. An example landing ellipse is overlaid on a terrain in Figure 5a. Reachable, marginally reachable, and unreachable points are shown in green, yellow, and red, respectively. A 20% margin of error is used to define the marginally reachable sites in Figure 5a (as well as the example results shown later).

V. SCIENCE FACTORS

Landing site selection for a space exploration mission is generally a compromise between safety and scientific return. When safety cannot be guaranteed, a potential site must be discarded—regardless of its potential scientific impact. Determining areas of high scientific potential is a laborious process that involves numerous considerations beyond the scope of an on-board reasoning system. It is, however, possible to integrate the scientists' preferred sites in order to influence the on-board site selection. Thus, for instance, the scientists may pre-select multiple potential sites that can be used in conjunction with the on-board terrain safety assessment in order to select the best site during descent. Such a scenario is considered here.

Assume scientists select a set of points of interest $(x_{0,i}, y_{0,i})$ in the terrain. Each point is at the center of a circular region of interest with radius r_i :

$$(x - x_{0,i})^2 + (y - y_{0,i})^2 \leq r_i^2 \quad (10)$$

The regions of scientific interest may or may not be ranked. If the sites are ranked, the ranking may be relative to other sites or based on a scale of interest. The general framework proposed here can account for any of these scenarios. In this paper, the ranking is based on level of interest. During the site selection process, scientists assign a score to each site based on the potential for scientific return. It is possible for multiple sites to have the same score.

At the point of entry, all pre-selected locations are reachable. As the terrain safety is assessed, the site that best combines safety, engineering, and scientific criteria is used for re-targeting. The process is done repeatedly during descent until arriving at a final selected landing site. Figure 5b shows an example terrain with three regions of interest, each with a particular science ranking. In this case, three different science levels (or rankings) are used: *high* (shown in green), *medium* (shown in yellow), and *low* (shown in orange). A fourth level represents zero scientific interest (shown in red).

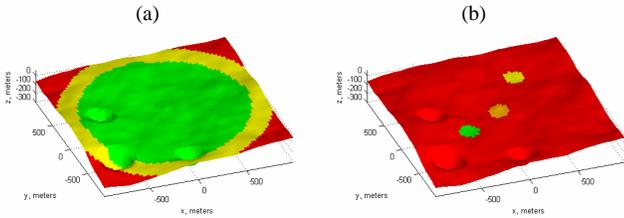


Fig. 5 Reachable (a) and scientifically interesting (b) regions of the terrain.

VI. LANDING SITE SELECTION

As discussed earlier, during spacecraft descent, the terrain safety is determined based on features extracted from the on-board sensors. Based on safety alone, a landing site could be selected by choosing the region of the terrain with the maximum a posteriori probability $P(T|\mathbf{f})$. However, if the spacecraft cannot reach the selected site or if it is of minimal

scientific interest, the selection is meaningless. Thus, engineering and science factors must be combined with terrain safety in order to select the landing site. Let L be a discrete random variable representing landing quality. The landing quality L depends on the state of the terrain T , the available fuel F , the and level of scientific interest S . Based on these causal relationships, a BN for landing site selection can be constructed, as in Figure 6.

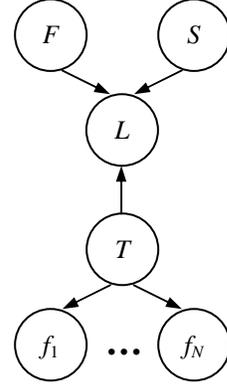


Fig. 6 BN for landing site selection.

Noting that nodes at the same level in a BN are conditionally independent, the joint probability encoded by the BN in Figure 6 can be written as:

$$P(T, F, S, L, f_1, \dots, f_N) = P(T)P(F)P(S)P(L|T, F, S) \prod_{i=1}^N P(f_i | T) \quad (11)$$

As can be seen, the bottom portion of the BN structure in Figure 6 is the same as the naïve Bayes structure for terrain safety.

For landing site selection, the quantity of interest is $P(L|T, F, S)$. This probability can be determined using causal (or *top-down*) reasoning. Let q represent the state of L at a point on the terrain. The quantity $P(L=q|T, F, S)$ represents the probability that the landing site has quality level q , given terrain safety, fuel sufficiency, and scientific interest. The expected quality of a potential site is

$$E(L|T, F, S) = \sum_q qP(L=q|T, F, S) \quad (12)$$

where $E(\cdot)$ is the expectation operator. The best landing site is the point on the terrain with the highest expected landing quality:

$$l^* = \max\{E(L|T, F, S)\} \quad (13)$$

The landing quality L is based on the possible states of T , F , and S . Thus, q can only take on a finite number of values. Setting $q \in [0, 1]$ ensures that $l^* \in [0, 1]$.

VII. EXPERIMENTS

A. Simulation

In order to evaluate the proposed approach, a series of experiments were performed using DSENDS, a high-fidelity dynamics and spacecraft simulator for entry, descent and landing [14]. Also, a suite of Digital Elevation Maps (DEMs) representing a variety of planetary terrains was created. The terrains are generated using a fractal model. Hills, craters, and rocks are added using appropriate models to ensure realism. Features are extracted from these DEMs using (RADAR, LIDAR and camera) sensor models while simulating spacecraft descent in DSENDS.

B. Supervised Learning

A supervised approach is used to learn the distributions in the Bayesian framework. Safety ground truth is obtained by estimating the final pose of the spacecraft at every point on the terrain DEM. In addition, rocks and craters are automatically deemed unsafe. A set of 10 different DEMs were used—each representing a different planetary landscape. Training for the terrain safety assessment and landing site selection is performed using the *leave-one-out* approach where nine DEMs are used for training and the remaining DEM is used as a test case. Each cell in the DEM is treated as an independent observation. It should also be noted that three different classifiers are trained, one for each descent tier (see Table I).

The features, \mathbf{f} , are a combination of discrete and continuous random variables. Specifically, the craters and rocks are discrete and the slope and roughness are continuous. A conditional Gaussian model is used and the likelihood $P(\mathbf{f}_i | T)$ is obtained using maximum likelihood estimation.

C. Results

Landing site selection results are shown in Figure 7. The safety assessment is overlaid on each terrain. The fuel ellipse is dashed and the sites of scientific interest are solid. Each site of scientific interest is centered about the original point selected by a scientist and shows a broad area with a circular radius of 100m. The science ranking is indicated by a corresponding *H* (high), *M* (medium), or *L* (low). The sites of scientific interest were not selected by actual scientists—they are only meant for evaluation purposes. The final selected landing site is shown with a black hash mark. In addition, the final landing score l^* is shown for each selected site. All examples are at an altitude of 4Km.

The results indicate that the proposed approach selects the site that best optimizes all three key factors: terrain safety, fuel, and science. For instance, in Figure 7d, there are two regions of high scientific interest. However, one of the regions lies near the boundary of the fuel ellipse and the other lies in an unsafe part of the terrain. Consequently, a site with medium scientific interest is selected because it has a better balance of all three key factors. In Figure 7b, the selected site does lie in a region of high scientific interest. Although this region straddles a crater, there are enough safe points from which to select a final landing site. A similar interplay between the three key factors is seen in the other examples.

V. CONCLUSIONS AND FUTURE WORK

This paper described a probabilistic approach to landing site selection during autonomous spacecraft descent. It was shown that BNs can adequately model a variety of criteria, including terrain safety, fuel consumption, and scientific interest, which can be used to determine the best landing site. Although the framework was proposed for autonomous landing, it could also be applied to navigation. Plans for future work include incorporating principles of probabilistic reasoning over time as well as dynamic sensor registration in order to allow for continuous reasoning with multiple re-targeting operations. In addition, the proposed method will be compared with other approaches in order to assess its relative performance. Considerable validation work will also be performed, including Monte Carlo simulations and experiments using real data collected from field experiments.

ACKNOWLEDGEMENTS

This work was supported by the Mars Technology Program, NASA Science Mission Directorate.

REFERENCES

- [1] M.P. Golombek, R.A. Cook, H.J. Moore, and T.J. Parker, "Selection of the Mars Pathfinder Landing Site," *Journal of Geophysical Research*, 102(E2), February 1997, pp.3567-3588.
- [2] M. Golombek et al., "Selection of the Final Four Landing Sites for the Mars Exploration Rovers," *Lunar and Planetary Science XXXIV*, 2003.
- [3] Y. Cheng, A. Johnson, C. Olson, L. Matthies, "Optical landmark detection for spacecraft navigation," *AAS/AIAA Space Flight Mechanics Meeting*, Ponce, Puerto Rico, February 2003.
- [4] Y. Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," *IEEE International Conference on Robotics and Automation*, April 2005.
- [5] B. Leroy, G. Medioni, A. E. Johnson, and L. Matthies, "Crater Detection for Autonomous Landing on Asteroids," *International Journal of Image and Vision Computing*, 19(11), September 2001.
- [6] A. Howard and H. Seraji, "Multi-Sensor Terrain Classification for Safe Spacecraft Landing," *IEEE Transactions on Aerospace and Electronic Systems*, 40(4), October 2004, pp.1122-1131.
- [7] A. Johnson, J. Montgomery, and L. Matthies, "Vision Guided Landing on an Autonomous Helicopter in Hazardous Terrain," *IEEE International Conference on Robotics and Automation*, April 2005.
- [8] A. Johnson, A. Klumpp, J. Collier, and A. Wolf, "LIDAR-Based Hazard Avoidance for Safe Landing on Mars," *AAS/AIAA Space Flight Mechanics Meeting*, Santa Barbara, CA, February 2001.
- [9] N. Serrano, M. Bajracharya, A. Howard, and H. Seraji, "A Novel Tiered Sensor Fusion Approach for Terrain Characterization and Safe Landing Assessment," *IEEE Aerospace Conference*, March 2006.
- [10] A. Huertas, Y. Cheng, and R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *IEEE Aerospace Conference*, March 2006.
- [11] J. Pearl, "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference," Morgan Kaufmann, 1988.
- [12] N. Friedman, D. Geiger, M. Goldszmidt, "Bayesian Network Classifiers," *Machine Learning*, 29, pp.131-163, 1997.
- [13] S.R. Ploen, H. Seraji, and C.E. Kinney, "Determination of spacecraft landing footprint for safe planetary landing," Submitted to *IEEE Transactions on Aerospace and Electronic Systems*.
- [14] J. Balaram, R. Austin, P. Banerjee, T. Bentley, D. Henriquez, B. Martin, E. McMahon, G. Sohl, "DSENDS - A High-Fidelity Dynamics and Spacecraft Simulator for Entry, Descent and Surface Landing," *IEEE Aerospace Conference*, March, 2002.

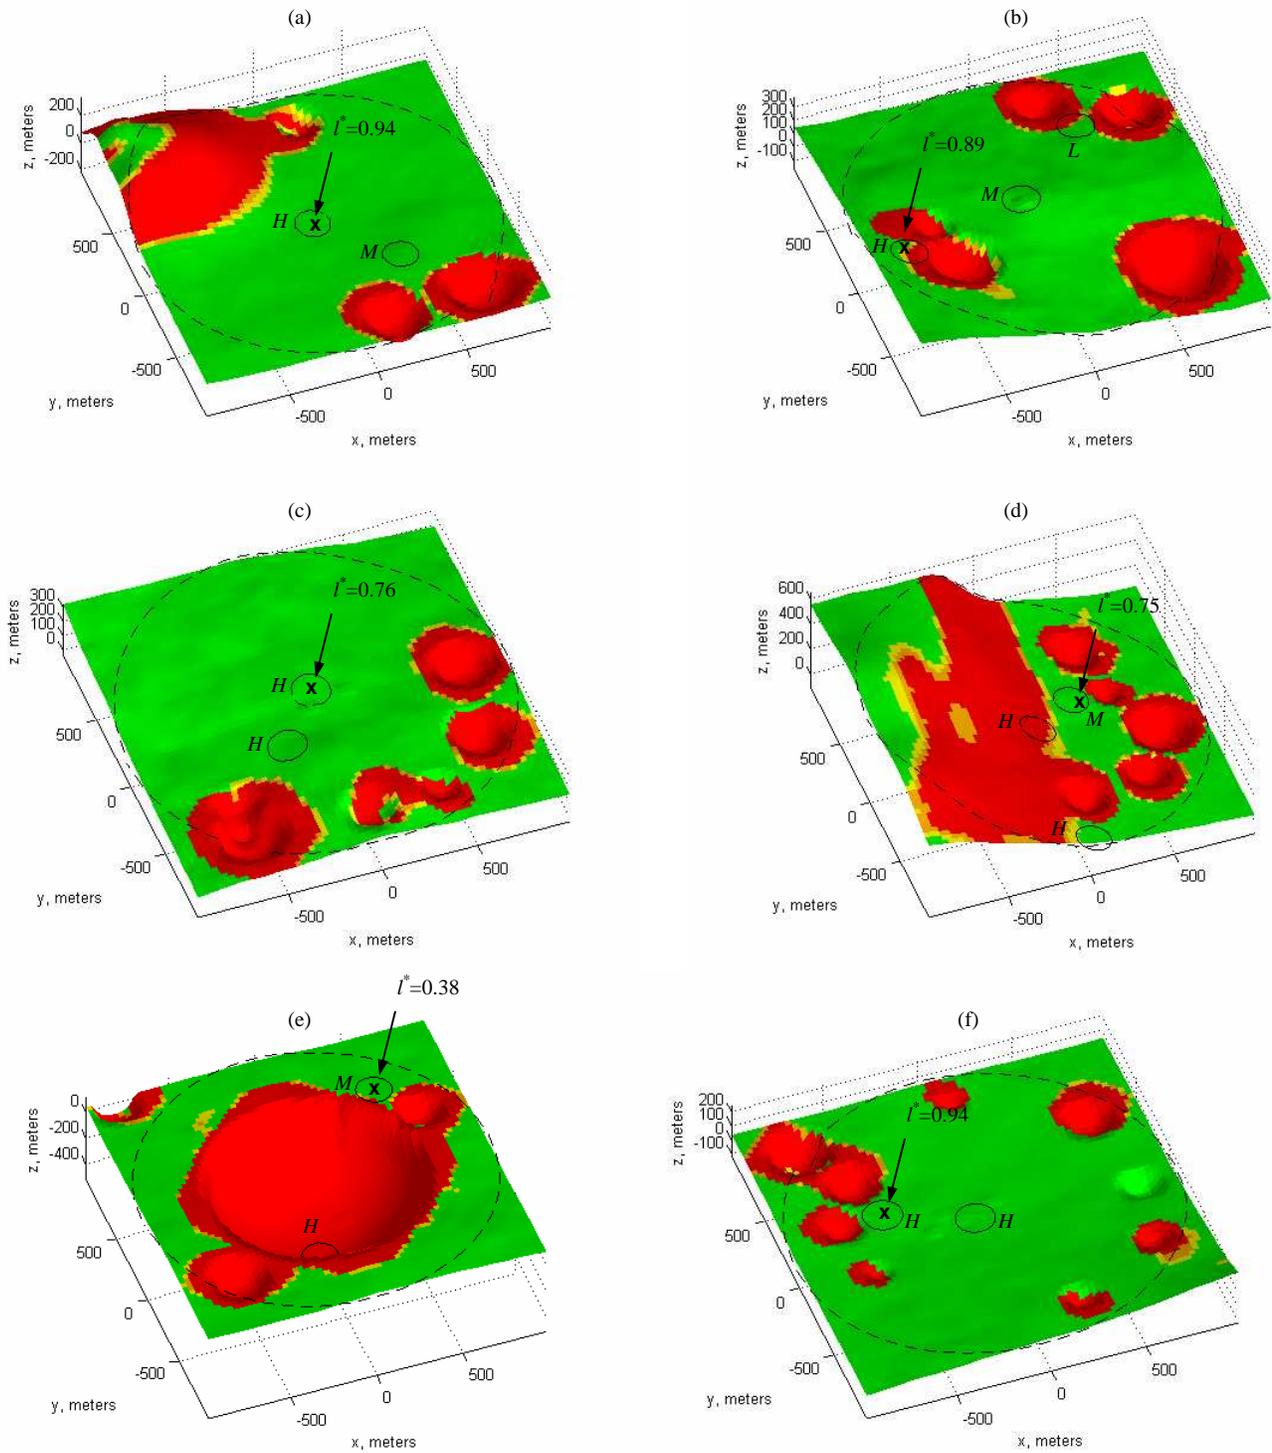


Fig. 7 Final landing sites selected for six different planetary terrains. The dashed ellipses represent the reachable area (based on fuel allocation) and the solid ellipses represent regions of scientific interest. The level of scientific interest is indicated by an *H* (high), *M* (medium) or *L* (low). The selected sites are shown with a hash mark. In addition, the landing score l^* for each of the selected sites is also shown.